

Robust likelihoods for inflationary gravitational waves from maps of cosmic microwave background polarization

Eric R. Switzer

NASA Goddard Space Flight Center, Greenbelt, Maryland 20771, USA

Duncan J. Watts

Department of Physics and Astronomy, The Johns Hopkins University,
3400 North Charles Street, Baltimore, Maryland 21218, USA

The B-mode polarization of the cosmic microwave background provides a unique window into tensor perturbations from inflationary gravitational waves. Survey effects complicate the estimation and description of the power spectrum on the largest angular scales. The pixel-space likelihood yields parameter distributions without the power spectrum as an intermediate step, but it does not have the large suite of tests available to power spectral methods. Searches for primordial B-modes must rigorously reject and rule out contamination. Many forms of contamination vary or are uncorrelated across epochs, frequencies, surveys, or other data treatment subsets. The cross power and the power spectrum of the difference of subset maps provide approaches to reject and isolate excess variance. We develop an analogous joint pixel-space likelihood. Contamination not modeled in the likelihood produces parameter-dependent bias and complicates the interpretation of the difference map. We describe a null test that consistently weights the difference map. Excess variance should either be explicitly modeled in the covariance or be removed through reprocessing the data.

I. INTRODUCTION

Systematic error control and rejection are central considerations of cosmic microwave background (CMB) instrument design and data analysis. The reionization feature of inflationary gravitational waves [25, 26, 43] and improved constraints on the optical depth τ [2] are of great interest, and require analysis of CMB polarization on the widest angular scales of the sky. These measurements are susceptible to contamination because they relate instrument response and foregrounds across the largest angular and temporal separations.

A well-established approach to detect or reject systematics splits the data into epochs, frequencies, surveys or other subsets across which a contaminant varies or is uncorrelated. This general approach can take the form of a cross-spectral estimator [51] or power spectrum null tests across a variety of difference maps. The maps can be split and subtracted to check for particular instrumental effects such as time constants (e.g. [3]) or to get uncorrelated realizations of detector noise or atmospheric fluctuations (e.g., [7, 44]). Contamination that is uncorrelated between the maps does not produce bias in the cross power, but it does boost errors. Here, we consider robustness and bias in the pixel-space likelihood and develop an approach analogous to the cross power or difference map null test.

Anisotropy spectral analysis compresses map information by exploiting the Gaussianity and statistical isotropy of the CMB signal. A survey is limited to fractions of the sky by its scan strategy and galactic contamination. Truncation has two consequences for the power spectrum. The estimate \hat{C}_ℓ is the sum of quadratic products of normally distributed map variations, making $P(C; \mathbf{x})$ (given map data vector \mathbf{x}) non-Gaussian unless there are sufficient modes available to be in the central limit. On a partial sky, the spherical harmonics for intensity and polarization are an incomplete basis, which results in correlations between ℓ [18, 24], and E- and B-mode

polarization mixing [5]. Polarization power spectra through pseudo- C_ℓ [24, 36] and related quadratic methods have well-established procedures for deriving the $P(C; \mathbf{x})$ [10, 16, 20]. Mangilli et al. [33] recently developed an implementation of a complete probability distribution function for the cross power spectrum on large scales, including the effects above.

An alternative approach resolves the challenges of spectral estimation by determining the cosmological parameters directly from the map [13, 28, 35, 38, 52]. The pixel-space likelihood approach has several advantages. Cuts and variation in coverage on a partial sky are included in the pixel covariance model and do not require simulations or analytic treatment of multipole correlation or polarization mixing. The likelihood represents the complete information contained in maps with Gaussian signal and noise, so it can achieve the lowest-variance estimates of the parameters. It accommodates some classes of foreground subtraction and self-consistently propagates parameter errors [28, 38, 52]. The likelihood avoids the need to represent a complete joint (non-Gaussian, correlated) probability distribution of the C_ℓ 's for ℓ on angular scales comparable to the survey size. Finally, the likelihood does not require reference to fiducial model parameters. There are well-established procedures for combining information from pixel-space likelihoods of large scales and spectral analysis of small scales [35, 38].

Drawing an analogy to the power spectrum null test, we can estimate cosmological parameters from the likelihood of a difference between two maps from subsets of the data. The sky signal drops out in the difference, and only noise or variable contamination remains. If there is no contamination, the posterior distribution of cosmological parameters will be consistent with data containing only detector noise. A null test based on the likelihood of the difference may fail to be informative when contamination does not match covariance in the data model. In this case, excess variance produces parameter-dependent bias through the effective weighting of the contam-

inants in map space. In the difference map, there is no signal, and the contaminant is weighted differently than in the sum.

We develop the joint likelihood analogy to the cross power and null test and propose a reweighting method for difference maps. Reweighting gives a consistent interpretation of bias from contamination in the sum and difference. However, excess variance (misspecification) also produces bias in the width of the posterior parameter distribution. Contamination must be treated either in a reanalysis of the data or as a new term in the likelihood model.

The methods developed here apply to experiments and missions specifically seeking the largest scales on the sky (Advanced ACTPol [21], CLASS [12], GroundBIRD [45], LSPE [1], PIPER [31], and QUIJOTE [32]) and missions (CORE+¹, Inflation Probe [30], LiteBIRD [34], and PIXIE [29]), but also to the largest angular scales of surveys on smaller regions.

Section II reviews the likelihood approach and its relation to spectral methods. Section III extends the likelihood to include the detector noise amplitude as a nuisance parameter. This setting demonstrates properties of an incomplete covariance model. Section IV reviews the cross-spectral estimator and defines an analogous joint likelihood between data splits. Section IV D describes the joint estimator with unknown contamination covariance and develops a re-weighted difference map null test. Section V summarizes the approach.

II. THE LIKELIHOOD OF CMB MAPS

The Gaussian log-likelihood $L \propto \ln P(x|\theta)$ for map data x given parameters θ is

$$2L = \text{Tr}[\ln C(\theta) + C(\theta)^{-1} D(x; \theta)] \quad (1)$$

$$D(x; \theta) \equiv [x - \langle x \rangle(\theta)][x - \langle x \rangle(\theta)]^T;$$

where C and $\langle x \rangle$ are the covariance and mean of the map vector x . In subsequent equations, we assume implicit parameter dependence in the mean and covariance model or emphasize dependence on individual model parameters.

The data vector for CMB polarization is $x^T \equiv (x_Q^T; x_U^T)$, a stack of Stokes Q and U maps. Stokes I intensity can also be added for complete TT , TE , EE , and BB two-point information, but models and simulations here use only Stokes Q and U for simplicity. The maps used in the likelihoods are smoothed at $\theta_{\text{FWHM}} = 15^\circ$ and binned onto $N_{\text{side}} = 8$ [15], encompassing multipoles $2 \leq l \leq 23$. The mask region is defined by WMAP P06 [35] and a declination limit of $-73^\circ < \delta < 27^\circ$ available to wide-area surveys in the Atacama, such as Advanced ACTPol [21] and CLASS [52] in the near term. Combined, $f_{\text{sky}} \approx 0.5$, shown in Fig. 1.

The total covariance $C(\theta)$ is the sum of cosmological signal and noise, $C(\theta) = S(\theta) + N(\theta)$. The signal covariance matrix is defined from the power spectrum $C_\ell(\theta)$ through [38]

$$S(\theta) = \sum_{\ell} \sum_{XY} C_\ell^{XY}(\theta) P^{XY}; \quad (2)$$

where our XY sum extends over only EE and BB to predict the Stokes Q and U covariance. P^{XY} is described in Tegmark & de Oliveira-Costa [48] and includes effects of the $\theta_{\text{FWHM}} = 15^\circ$ smoothing and incomplete sky coverage. For simplicity, we fix Λ CDM cosmological parameters [23] throughout, but these could be jointly estimated in the likelihood. Conclusions are not dependent on the base cosmology at the currently available precision. The input data x are generated using *synfast* [15] plus a Gaussian random detector noise and are consistent with the covariance in Eq. (2). In subsequent sections, we add contamination to the maps to study departures from the model.

For the purpose of demonstrating the B-mode constraint, we use a reference survey with noise RMS amplitude $10 \mu\text{K arcmin}$. This noise level is typical of per-band sensitivities of next-generation experiments in their prime CMB science band (see e.g. Errard et al. [11] and Remazeilles et al. [39] for summaries). Throughout, “detector noise” refers to the variance in the map attributable to the detectors, which we assume is uncorrelated between map pixels.

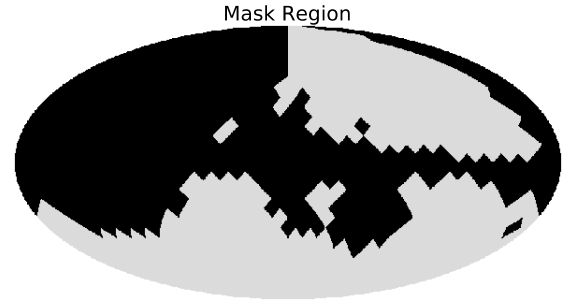


FIG. 1. Mask used in the simulations here ($N_{\text{side}} = 8$, giving $f_{\text{sky}} \approx 0.5$), representative of the region accessible from the Atacama, as well as masking the galaxy through WMAP’s P06 map.

To simulate the distribution of experimental outcomes, we find maximum likelihood (ML) parameters across Monte Carlo realizations of data sets using the limited-memory Broyden-Fletcher-Goldfarb-Shanno algorithm [6, 53].

A. Relation to quadratic methods

The Newton-Raphson approach provides an analytic expression to iterate to find the maximum likelihood as [4, 9]

$$\delta\theta = \sum_j \frac{1}{F_{ij}} \frac{1}{2} \text{Tr} \left(x x^T - C \right) C^{-1} C_{ij} C^{-1} \delta\theta_j \quad (3)$$

$$F_{ij} \equiv \frac{1}{2} \text{Tr} (C^{-1} C_{ij} C^{-1} C_{ij}); \quad (4)$$

where commas denote partial derivatives as $C_{,i} \equiv \partial C / \partial \theta_i$. This expression gives the update to the parameter θ_i among the parameters in $\theta = \{\theta_1, \dots, \theta_g\}$. Equation (3) evaluates all covariances and derivatives at the current iteration, and the resulting $\delta\theta$ gives the vector of changes $\delta\theta$ to iterate to the maximum

¹

<http://conservancy.umn.edu/handle/11299/169642>

likelihood. Each iteration is driven by $\mathbf{x}\mathbf{x}^T \square \mathbf{C}(\square)$, the difference between the covariance model at that iteration and the outer product of the data.

The covariance model $\mathbf{C} = \square \mathbf{C}_\square + \mathbf{N}$ provides a simple example of quadratic estimation of the covariance amplitude \square . Let the first guess be $\square = 0$. Then the first step toward the maximum likelihood value of \square is

$$\hat{\square} = \frac{\mathbf{x}^T \mathbf{Q}_\square \mathbf{x} \square \mathbf{b}}{\text{Tr}(\mathbf{Q}_\square \mathbf{C}_\square)} \quad \mathbf{b} = \text{Tr}(\mathbf{Q}_\square \mathbf{N}) \quad (5)$$

$$\mathbf{Q}_\square = \mathbf{N}^{-1} \mathbf{C}_\square^{-1} \mathbf{N}^{-1} \quad (6)$$

The quadratic term $\mathbf{x}^T \mathbf{N}^{-1} \mathbf{C}_\square^{-1} \mathbf{N}^{-1} \mathbf{x}$ inverse-noise weights the map data ($\mathbf{N}^{-1} \mathbf{x}$) and dots across the covariance structure \mathbf{C}_\square . The “noise bias” \mathbf{b} removes the contribution of noise \mathbf{N} . The denominator is a normalization that ensures that the expectation value is

$$\hat{\square} = \frac{\text{Tr}(\square \mathbf{Q}_\square \mathbf{C}_\square + \mathbf{Q}_\square \mathbf{N}) \square \text{Tr}(\mathbf{Q}_\square \mathbf{N})}{\text{Tr}(\mathbf{Q}_\square \mathbf{C}_\square)} = \square \quad (7)$$

The quadratic estimator with full \mathbf{C}^{-1} weights is the minimum-variance estimate of Gaussian covariance amplitude parameters [47, 48]. Quadratic methods are commonly used to estimate the anisotropy spectrum [8–10, 14, 19, 24, 40]. In this case, the parameters of the covariance model are the \mathbf{C}_\square ’s themselves or band powers.

The quadratic approach to the maximum likelihood provides some analytic intuition about the behavior of the maximum likelihood (Sec. III A) and relation of the joint pixel-space likelihood to the cross power (Sec. IV A and Appendix A).

B. Applicability of the likelihood approach

The pixel-space likelihood is analytically simple at the expense of being computationally intensive and structurally rigid. Evaluation of the likelihood requires the specification and inverse of the $\mathbf{N}_{\text{pix}} \square \mathbf{N}_{\text{pix}}$ covariance matrix of the maps, which is numerically expensive. The pixel-space likelihood has therefore seen greatest use in extracting information from the largest scales in the survey, which span $\mathcal{O}(1000)$ pixels. Note that the pixel-space analysis is generally useful for modeling signal covariance on angular scales approaching the size of the survey, not just at low- l .

Given map data and a model, the likelihood is a self-sufficient “black box” to determine the cosmological parameters. As long as the data model is accurate, the likelihood gives the probability distribution of the parameters. Spectral methods provide greater freedom, such as choosing spatial weightings, removing foregrounds in advance, and throwing out spatial modes. Freedom also carries the responsibility of propagating treatments to the final parameter distribution. Choices in data weighting and filtering other than \mathbf{C}^{-1} move away from optimality. In contrast, the likelihood has only the fixed model and does not differentiate between the signal structure covariance and data weighting or filtering.

Covariance in the observed data that is neither modeled nor isolated could produce a spurious detection of B-modes or

bias the true value. Some contaminant covariance structure may be known accurately in advance. For example, the variation of detector noise across the map can be modeled from the survey coverage and data cuts. However, the amplitude of that noise may not be known from only laboratory measurements or characterization of the time domain data. It is straightforward to include detector noise amplitude in the likelihood’s covariance model and to fit it jointly with the cosmological parameters. Fitting for the amplitude removes that source of bias from the cosmological parameters and gives confidence regions that reflect the full covariance.

III. BIAS FROM EXCESS DETECTOR NOISE

A simple covariance model for CMB polarization data is the sum of tensor and scalar cosmological contributions and a known noise, or [28, 52]

$$\mathbf{C}(\mathbf{r}; \mathbf{s}; \mathbf{g}) = \mathbf{r} \mathbf{S}_T + \mathbf{s} \mathbf{S}_S + \mathbf{N}; \quad (8)$$

where the tensor-to-scalar ratio \mathbf{r} multiplies the tensor covariance structure \mathbf{S}_T , the scalar amplitude $\mathbf{s} \square \mathbf{A}_s = \mathbf{A}_{s,0}$ multiplies scalar covariance \mathbf{S}_S , and \mathbf{N} is fixed detector noise. \mathbf{S}_T and \mathbf{S}_S are derived from the sum on EE and BB in Eq. (2) for fixed \square CDM cosmological parameters [23].

Figure 2 shows the Monte Carlo maximum likelihood distribution of \mathbf{r} for an input $\mathbf{r} = 0.05$ when both the modeled noise and the true map noise are $10 \square \text{K arcmin}$, and also the case when the detector noise is 30% higher (in map space) than modeled in the fixed \mathbf{N} term.

A simple extension to the covariance model can also fit for the detector noise amplitude in the maps with the covariance model, as [17]

$$\mathbf{C}(\mathbf{r}; \mathbf{s}; \square \mathbf{g}) = \mathbf{r} \mathbf{S}_T + \mathbf{s} \mathbf{S}_S + (\square = \square_0)^2 \mathbf{N}; \quad (9)$$

Figure 2 shows that the distribution of \mathbf{r} is correctly centered around the input $\mathbf{r} = 0.05$ when the noise amplitude of $13 \square \text{K arcmin}$ is jointly modeled with the cosmological signal as in Eq. (9). It is slightly broader due to the higher level of detector noise and expense of fitting \square .

It is possible to fit for cosmological amplitudes \mathbf{r} and \mathbf{s} in parallel with noise \square because these have different covariance structures. In multipole space, \mathbf{s} modulates the E-modes, \mathbf{r} impacts the E- and B-modes and detector noise contributes to both but with a different l dependence. When a contaminant has more overlap with the B-mode covariance structure, it becomes harder to separate. In the worst case, the excess covariance structure is identical to the signal and so is indistinguishable.

Figure 3 shows the $1\square$ and $2\square$ equivalent regions for the distribution of bias in $\hat{\mathbf{r}}_{\text{ML}}$ that is produced by detector noise 30% higher than modeled. The bias has clear dependence on the amplitude of true B-modes present in the map. The parameter dependence of the bias goes against intuition from the power spectrum, where uncorrelated contributions in the map are additive in \mathbf{C} . The behavior in Fig. 3 does not occur for all types of covariance. For example, adding $\mathbf{r} = 0.1$ B-modes to a map with $\mathbf{r} = 0$ B-mode amplitude will produce

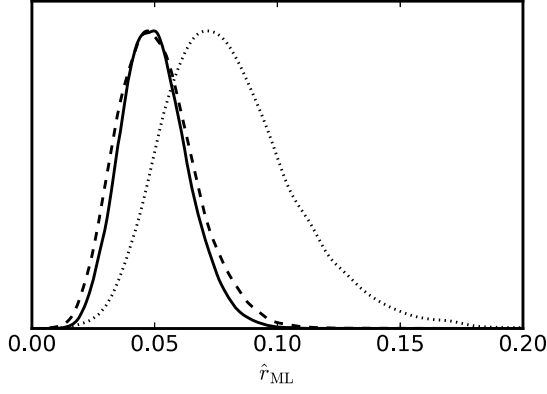


FIG. 2. The Monte Carlo distribution of the maximum likelihood tensor-to-scalar ratio r for several scenarios of detector noise. Solid curve: the distribution of r when both the map and the covariance model have 10 K arcmin noise, and a true input $r = 0.05$. Dotted curve: the biased distribution of r when the map noise is 13 K arcmin while the likelihood's covariance model assumes 10 K arcmin . Dashed curve: the distribution of r with a map noise level of 13 K arcmin , and the likelihood models the detector noise amplitude jointly with the signal. Jointly fitting for the noise recovers the input r value and represents the larger uncertainty in r from 13 K arcmin noise vs. the 10 K arcmin reference case.

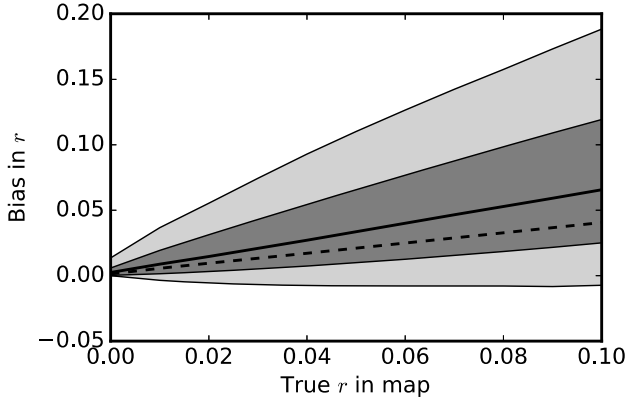


FIG. 3. The bias in r_{ML} from detector noise that is 30% higher than modeled in the likelihood, as a function of the true B-mode r signal in the map. Light and dark gray regions are percentile regions equivalent to 2σ and 1σ respectively, about the median (solid line). The distribution is estimated as the Monte Carlo of the maximum likelihood value of r over map realizations of both signal and detector noise. The dashed line shows bias in r inferred from the quadratic estimate of r , as the first step of the Newton-Raphson approach to the maximum likelihood (Sec. II A). The presence of true B-mode signal modulates the weighting of variance that contributes to the bias.

an estimate of $\langle r \rangle = 0.1$. Adding a systematic with equivalent $r = 0.1$ B-modes to a map with true signal $r = 0.1$ B-mode amplitude produces an estimate of $\langle r \rangle = 0.2$. More generally, for a systematic level of r_{sys} , the measured $\langle r \rangle = r_{\text{true}} + r_{\text{sys}}$, independently of the level of r_{true} .

A. Parameter dependence of the bias

There is no intuitive closed-form expression for the maximum likelihood, but the quadratic approach to the maximum likelihood can give approximate expressions that explain the basic behavior. The quadratic estimate for \hat{r} in Eq. (8) is analogous to the covariance amplitude determination in Eq. (5), and gives

$$\hat{r} = (\mathbf{x}^T \mathbf{C}^{-1} \mathbf{S}_T \mathbf{C}^{-1} \mathbf{x} - \mathbf{b}) / \text{Tr}(\mathbf{C}^{-1} \mathbf{S}_T \mathbf{C}^{-1} \mathbf{S}_T) \quad (10)$$

$$\mathbf{b} = \text{Tr}(\mathbf{C}^{-1} \mathbf{S}_T \mathbf{C}^{-1} \mathbf{N}); \quad (11)$$

where matrices are as defined in Eq. (8). Taking the expectation value gives $\text{Tr}(\mathbf{C}^{-1} \mathbf{S}_T \mathbf{C}^{-1} \langle \mathbf{h} \mathbf{x} \mathbf{x}^T \mathbf{i})$ in the numerator. Identify the total covariance

$$\mathbf{C}_{\text{tot}} \equiv \langle \mathbf{h} \mathbf{x} \mathbf{x}^T \mathbf{i} = r \mathbf{S}_T + s \mathbf{S}_S + \mathbf{N} + \mathbf{\Omega}; \quad (12)$$

where $\mathbf{\Omega}$ is some contamination covariance present in the data but not the model. Thermal noise is represented by \mathbf{N} in the model, and is removed through subtraction of \mathbf{b} . The remaining incurred bias in \hat{r} due to the un-modeled term $\mathbf{\Omega}$ is

$$\langle \hat{r} \rangle_{\text{bias}} = \text{Tr}(\mathbf{C}^{-1} \mathbf{S}_T \mathbf{C}^{-1} \mathbf{\Omega}) / \text{Tr}(\mathbf{C}^{-1} \mathbf{S}_T \mathbf{C}^{-1} \mathbf{S}_T); \quad (13)$$

In the numerator, $\mathbf{C}^{-1} \mathbf{\Omega}$ weighs the contamination $\mathbf{\Omega}$, and \mathbf{C} depends on the cosmological parameters r and s through covariance terms $r \mathbf{S}_T + s \mathbf{S}_S$. Hence the weighting of the contamination is parameter dependent. Other parameters such as $\langle \hat{r} \rangle_{\text{bias}}$ have a similar expression of some weighted overlap of the signal covariance structure with the contamination. The contamination biases all parameters that have structural overlap, and in a way that depends on the value of the parameters. Figure 3 shows the quadratic approximation to the bias as a dashed line. It is lower than the median of maximum likelihood realizations but captures much of the effect. Recall that the quadratic expression is only the first step of the Newton-Raphson approach to the maximum likelihood.

To get additional analytic intuition, Appendix B derives $\langle \hat{r} \rangle_{\text{bias}}$ in terms of the signal and contaminant covariance eigenvalues in the case where the contaminant is a multiplier times the identity matrix, $\mathbf{\Omega} = \Omega_b^2 \mathbf{I}$. Then the bias becomes a simple ratio

$$\langle \hat{r} \rangle_{\text{bias}} = \Omega_b^2 \frac{\mathbf{w}(r)^T \mathbf{1}}{\mathbf{w}(r)^T \mathbf{\Omega}}; \quad (14)$$

where $\mathbf{w}(r)$ is the weight per signal eigenmode (given in Appendix B), and $\mathbf{\Omega}$ is the vector of signal eigenvalues. Figure 4 shows the terms of $\mathbf{w}(r)^T \mathbf{\Omega}$ for each mode, as a function of r , and is $\mathbf{S} = (\mathbf{S} + \mathbf{N})$ per mode. For low r , there are fewer signal-dominated modes. The r -dependence of the weight gives r -dependence to the bias. If the contaminant has the same form as the B-mode, then the numerator becomes $\mathbf{w}(r)^T \mathbf{\Omega}$, and the bias is independent of any true r in the map.

The quadratic estimator in Eq. (5) has additional flexibility that is not available in the likelihood because it separates the covariance structure of the signal from the covariance structure of the weights. That is, in the quadratic form

$$\hat{\mathbf{r}} / \mathbf{x}^T \tilde{\mathbf{C}}^{-1} \mathbf{C}_r \tilde{\mathbf{C}}^{-1} \mathbf{x}; \quad (15)$$

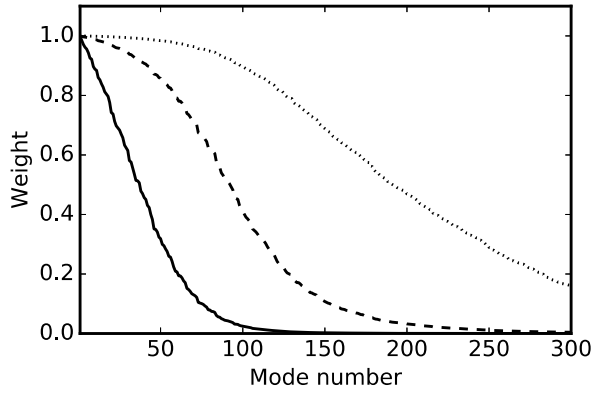


FIG. 4. $S/(S + N)$ weight per eigenmode for three values of tensor amplitude: $r = 0.01$ (solid curve), $r = 0.1$ (dashed curve), and $r = 1$ (dotted curve). The parameter dependence of the weight produces parameter dependence in the bias from contamination.

\tilde{C}^{-1} and C_{\square} do not need to originate from the same parent C . In the likelihood, this choice is fixed for both. Different choices of weights affect the optimality of the estimator and the structure of parameter correlations. In the pseudo- C approach to spectral estimation, the quadratic combination of data can be weighted by some N^{-1} (hit map) or other apodization W . In either case, the weighting does not depend on the parameters.

IV. LIKELIHOODS ACROSS DATA SPLITS

Many sources of contamination vary or become uncorrelated across subsets of data from different epochs, frequencies, surveys, or data treatments. The cross power estimator [51] extends the quadratic estimator in Eq. (5) by forming the quadratic product across subsets A and B of the data as

$$\hat{C} = x_A^T N_A^{-1} S_{\square} N_B^{-1} x_B. \quad (16)$$

The expectation value $\langle \hat{C} \rangle = \text{Tr}(N_A^{-1} S_{\square} N_B^{-1} \langle x_B x_A^T \rangle)$ contains $x_B x_A^T$, which averages to zero for any variance terms that are not common to both A and B. (Variance not common between A and B increases the variance of the estimator.) This approach has recently been extended to analysis of variance when few modes are available in a survey volume [33].

The difference of maps across a data split, $x_A - x_B$, will remove any astronomical signal common to both maps. The power spectrum of the difference map tests for any excess variance. For example, the analysis of time-domain data must account or compensate for the detector response time constants. Otherwise, time constants can produce a residual variance in the difference between maps of left- and right-going scans. Large suites of such null tests support the ultimate parameter determination by ruling out sources of contamination. This approach has been applied extensively to spectral analysis and, to a much more limited degree, pixel-space likelihoods [38].

A. The joint likelihood

The pixel-space likelihood analogy to the cross power is the joint likelihood of the maps A and B in the data split. Model the joint likelihood across the data split between A and B, with $x^T = (x_A^T \ x_B^T)$ as

$$C = \begin{pmatrix} S(\square) + N_A & S(\square) \\ S(\square) & S(\square) + N_B \end{pmatrix}; \quad (17)$$

where $S(\square)$ is in common to both A and B. Recall that x_A is still a stack of the Stokes Q and U maps, so the combined data vector $(x_A^T \ x_B^T)$ is the stack of four maps and the covariance is also naturally 4×4 blocks for correlations of Stokes Q and U across A and B. The noise covariance can also be extended to accept parameters such as the \square amplitude in Sec. III.

A likelihood model for that data that uses Eq. (17) extracts parameter information from $A \square B$ (cross) but also $A \square A$ and $B \square B$ (auto). This can be seen in the form of the quadratic estimator $x^T C^{-1} C_{\square} C^{-1} x$, where both C^{-1} and C_{\square} have off- and on-diagonal block terms, so the inner product with $x^T = (x_A^T \ x_B^T)$ mixes both $A \square B$ and $A \square A$ or $B \square B$ information. A likelihood that uses the covariance model of Eq. (17) will have sensitivity to information in $A \square A$, and so does not have the same immunity to uncorrelated noise as the cross power in Eq. (16).

To reach a closer analog to the cross power, take a covariance model which has a duplicate set of nuisance parameters on the diagonal as

$$C = \begin{pmatrix} S(\square) + S(\square_b) + N_A & S(\square) \\ S(\square) & S(\square) + S(\square_b) + N_B \end{pmatrix}; \quad (18)$$

Adding $S(\square_b)$ to the diagonal and marginalizing over \square_b effectively sweeps the rug out from under parameter constraints on \square coming from $A \square A$ and $B \square B$. Appendix A relates Eq. (18) to the cross quadratic product $x_A^T N_A^{-1} C_{\square} N_B^{-1} x_B$ (with no noise bias to remove) as the first step of a Newton-Raphson iteration starting from a noise-only covariance.

The approach in Eq. (18) resembles mode avoidance or cleaning strategies. These are commonly implemented by fitting and subtracting mode functions. Several authors [41, 42, 46, 49] consider this class of avoidance and argue that the following three are equivalent: (1) least-squared fitting and subtracting modes in the mean model, (2) marginalizing over the amplitude of nuisance modes in the mean model, and (3) taking a multiplier of the covariance structure of the contaminated modes to infinity. In the present case, rather than marginalizing over the bias nuisance variables \square_b , another approach that suggests itself would be to set the amplitude of these on-diagonal variance terms to infinity through \square_b . This limit throws away dependence on the nuisance parameters, and so halves the number of free parameters to estimate through the likelihood. Taking on-diagonal signal variance to infinity results in a quadratic estimator that uses only information across the data split, as intended. However, it also results in a C^{-1} weight applied to each map with infinite variance in the signal modes; e.g., it also eliminates the signal. Rather than marginalize over nuisance parameters in the mean, the

likelihood using the covariance model in Eq. (18) marginalizes over nuisance parameters in the covariance. Taylor & Kitching [46] give analytic expressions for marginalization over covariance parameters. These can be used to simplify computation in a high-dimensionality parameter space, but are not needed for the simple, few-parameter models considered here.

B. The sum-difference likelihood

Rotating to a basis of sum and difference maps, $(x_s, x_d) = (x_A^T + x_B^T, x_A^T - x_B^T)$ simplifies the joint covariance in Eq. (18). Take the detector noise amplitude to be identical between splits A and B for simplicity (this can be arranged by splitting complementary sets of the data with common integration depth). The resulting likelihood separates as the product of sum and difference likelihoods

$$\begin{aligned} P(x_j; \square_b) &= P_s(x_{sj}; \square_b) P_d(x_{dj}; \square_b) \\ P_s(x_{sj}; \square_b) &\propto N(0; 4S(\square) + 2S(\square_b) + 2N) \\ P_d(x_{dj}; \square_b) &\propto N(0; 2S(\square_b) + 2N): \end{aligned} \quad (19)$$

The factors of 2 and 4 are a by-product of taking the $x_A^T + x_B^T$ combination rather than $1/2(x_A^T + x_B^T)$. The signal in common to both maps appears as $4S(\square)$ because it adds coherently between the two maps, as $(2\square)^2$, and the bias of uncorrelated contamination appears in both the sum and the difference as $2S(\square_b)$ because it is the addition of two uncorrelated variances $S(\square_b) + S(\square_b)$ in each map.

An interpretation of the joint sum-difference likelihood of Eq. (19) is that P_s constrains the parameters plus bias, while the difference P_d constrains only bias from modulated contamination.

C. Sum-difference likelihood: B-mode contamination

Following Sec. III, take a simple model where $S(\square) = r_{\text{true}} S_T + s_{\text{true}} S_S$. Again S_T and S_S are the covariance structure of the tensor and scalar modes in the Stokes Q and U maps. The new parameters in the bias space are $S(\square_b) = r_{\text{bias}} S_T + s_{\text{bias}} S_S$. Figure 5 shows the sum and difference likelihoods for r for a simulation with $r_{\text{true}} = 0.05$, detector noise of $10 \mu\text{K arcmin}$ in each map, and contamination at the level of $r_{\text{bias}} = 0.1$ in uncorrelated realizations added to both maps. The sum map can only constrain $2r_{\text{true}} + r_{\text{bias}}$ so it is a degenerate band from upper left to lower right. The difference map can only constrain r_{bias} with no dependence on r_{true} , so it appears as a horizontal band. The product of the two likelihoods recovers both the input $r_{\text{true}} = 0.05$ and the contamination level. The factor of two in $2r_{\text{true}} + r_{\text{bias}}$ is effectively a normalization for the bias amplitude under the assumption that the source of bias is uncorrelated between the split maps.

Marginalizing over r_{bias} in the joint likelihood gives an estimate of r_{true} with errors self-consistently inflated to reflect the fact that some of the constraining power of the map is used to estimate bias. An analogous power-spectral null test uses the difference of maps to rule out bias parameters. If the power

spectrum of the difference map is consistent with zero, a particular source of contamination can be ruled out. In a typical use of the null test, once a source of variance is ruled out, it is taken to be identically zero. In contrast, in the joint likelihood, Eqs. (19) and (18), uncertainty in the bias parameter is folded into the final estimate \hat{r}_{true} .

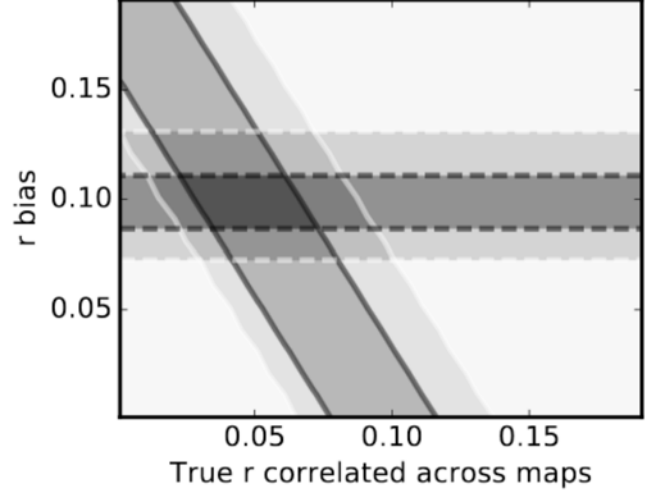


FIG. 5. The likelihood of r_{bias} and r_{true} from the sum and difference of maps across a split in the data. Here $r_{\text{true}} = 0.05$ is in common between the maps, and each map has an independent realization of contaminating B-modes at the level $r_{\text{bias}} = 0.1$. This represents a scenario where there is time-varying contamination with structure indistinguishable from the B-modes. Solid lines: The regions of 68% and 95% probability enclosed (1σ , 2σ equivalent) in the likelihood of the sum map. It traces a degeneracy between interpretation of the data as all true B-mode signal or all bias. Dashed lines: The analogous regions for the likelihood of the difference map. This likelihood isolates residual variance at the level $r_{\text{bias}} = 0.1$. The joint likelihood of sum and difference maps determines both the bias and the target r amplitude with self-consistent errors. Both posterior distributions are from a single realization of data, so the maximum likelihood is not centered on the input parameters.

When the likelihood model includes all covariance terms, the tensor to scalar ratio can be recovered without bias and with correct confidence intervals. When the contaminant covariance differs in structure from terms in the covariance model, the cross likelihood fails and can give biased results due to the weighting effect in Sec. III A.

D. Reweighting for difference map null tests

Section III provides a scenario where the data have some covariance that is not explained by the structure in the likelihood model. In the case of a single map, Sec. III A showed how a parameter \square is biased when its structure C_{\square} overlaps with contamination \square through $\text{Tr}(C_{\square}^{-1} C_{\square} C_{\square}^{-1} \square)$. Further, the bias depends on parameters through the $C(\square)^{-1}$ weight. In the sum-difference formalism, the difference map has no signal by construction, so the contamination is weighted differently than in the sum map likelihood.

Mis-specification of the covariance model results in significantly different parameter biases in the likelihoods of the sum and difference maps. In the difference map, there is no cosmological B-mode signal by construction, so the pixel-space covariance is equivalent to the case of $r = 0$. Figure 3 shows the bias produced by a 30% excess of thermal noise as a function of r assumed in the covariance. At $r = 0$, the likelihood of the difference map reports negligible bias produced by the excess thermal noise. In contrast, if the data have a true $r = 0.1$, the 30% excess thermal noise will produce a bias of $\Delta r = 0.06$ on average, with fluctuations at the level of $\Delta r = 0.05$. The parameters inferred from the likelihood of the difference map no longer provide useful information about the bias of parameters in the sum map, and may lead to false confidence in the analysis.

For the likelihood of the difference to constitute a null test, it must weigh the contamination consistently with the sum. An approach to consistent weighting is to add a signal realization x_{sig} to the difference map x_d and find the Monte Carlo average of the parameters over signal realizations (each realization will have some signal variance). Any deviation from the input parameters could be attributed to contamination in the difference map. In this case, the data matrix $\mathbf{h} = \mathbf{h}(x_{\text{sig}} + x_d)(x_{\text{sig}} + x_d)^T = x_d x_d^T + 4\mathbf{S}(\square)$, where \mathbf{h} is the expectation over signal realizations. (In the sum map, two coherent copies of the signal are added, giving a factor of 4 in variance.) Rather than Monte Carlo, we take the expectation value of the likelihood over added signal, giving the model

$$2L_d = \text{Tr}[\ln \tilde{\mathbf{C}}_d + \tilde{\mathbf{C}}_d^{-1}(x_d x_d^T + 4\mathbf{S}(\square))] \quad (20)$$

$$\tilde{\mathbf{C}}_d = 2\mathbf{S}(\square_b) + 2\mathbf{N} + 4\mathbf{S}(\square):$$

The factor of 2 in the bias covariance $2\mathbf{S}(\square_b)$ represents the assumption that the excess variance producing the systematic is not correlated across the difference of maps.

The role of $4\mathbf{S}(\square)$ in the covariance model is clear as a reweighting, but the $4\mathbf{S}(\square)$ term in the data matrix also plays an important role in the likelihood. In the Newton-Raphson approach to maximum likelihood, each iteration is based on the difference between the data matrix (outer product of the data) and the covariance model,

$$\square = \mathbf{D} - \mathbf{C}(\square) \quad (21)$$

$$= [x_d x_d^T + 4\mathbf{S}(\square)] - [2\mathbf{S}(\square_b) + 2\mathbf{N} + 4\mathbf{S}(\square)] \quad (22)$$

$$= x_d x_d^T - 2\mathbf{S}(\square_b) - 2\mathbf{N}: \quad (23)$$

The maximum likelihood therefore fits the residual variance in the difference map to the signal bias model, accounting for thermal noise. Recall that each Newton-Raphson step in Eq. (3) is weighted by a \mathbf{C}^{-1} , which also contains $4\mathbf{S}(\square)$ and weights \square consistently with the sum map.

The reweighted likelihood of the difference map should be interpreted as $P_d(\square_b | \square)$, the distribution of bias parameters evaluated in a map where signal variance is fixed at \square . It should not be interpreted as the joint likelihood $P_d(\square_b, \square)$.

The sum and difference likelihoods can be sampled and combined through the following process:

1. Use the map sum $x_A + x_B$ to constrain $2r_{\text{true}} + r_{\text{bias}}$, marginalized over all other parameters. This gives a diagonal band of degeneracy in the $r_{\text{true}}-r_{\text{bias}}$ plane.
2. Use the difference map $x_A - x_B$ in the re-weighted likelihood Eq. (20) (r_{true} fixed) to estimate the distribution of r_{bias} , marginalized over all other parameters. This gives a slice of the probability of r_{bias} for a given true level of signal r_{true} , $P(r_{\text{bias}} | r_{\text{true}})$.
3. Repeat the difference analysis for any other null test combinations, giving contours in the $r_{\text{true}}-r_{\text{bias}}$ plane.

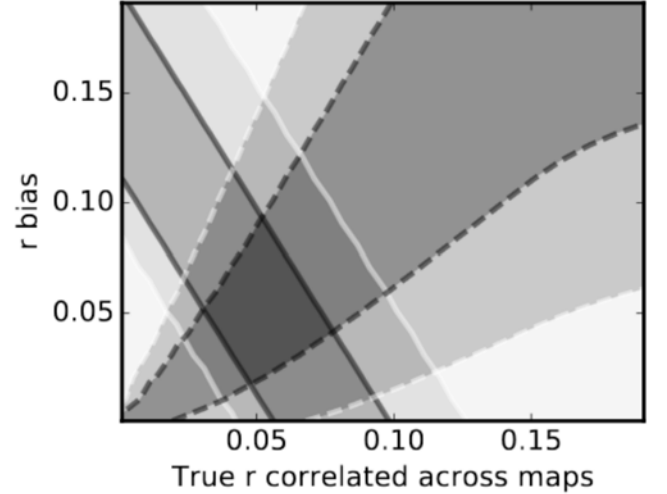


FIG. 6. Same layout as Fig. 5, except that instead of adding variance in the form of B-modes (which are in the likelihood covariance model), this simulation has detector noise 30% higher than modeled and is not explained by any free term of the covariance. Mis-specification produces a parameter-dependent bias. Here we force the likelihood of the difference map to weight consistently with the covariance in the sum map likelihood. Without reweighting, the likelihood of the difference map gives $r_{\text{bias}} < 0.004$ at 95% confidence.

Figure 6 applies this process to maps with $r_{\text{true}} = 0.05$ and detector noise that is 30% higher than is modeled in the likelihood. Without reweighting, the likelihood of the difference map reports $r_{\text{bias}} < 0.004$ at 95%, independently of r_{true} . Reweighting the likelihood results in $P(r_{\text{bias}} | r_{\text{true}})$ that depends on r_{true} , analogously to Fig. 3.

WMAP [22] and Planck [38] find the posterior distribution of \square from difference maps, but do not describe added signal covariance in the model. This covariance is required to consistently weight contamination in the null analysis.

To constitute a useful null test, the likelihood of the difference map must also give an informative confidence interval. Appendix C describes the curvature of the posterior parameter distribution. When there is no excess variance, the curvature is the usual Fisher matrix $\text{Tr}[\mathbf{C}^{-1} \mathbf{C}_{;\square} \mathbf{C}^{-1} \mathbf{C}_{;\square}]$. The width coincides with the distribution of the cosmological signal and noise at fixed contamination. The width is analogous to a standard null test, where the errors in the difference map power spectrum do not account for any contamination.

The width of the likelihood is erroneous when the difference map has an excess variance that is not described by free parameters in the likelihood model. Contamination must be treated either in a reprocessing of the data or be modeled self-consistently in the likelihood. Examples here could include a fit for noise amplitude in the map or deprojection [3], where instrumental systematics of known structure produce correlations between temperature and polarization maps. Extensions to the likelihood must balance adequacy (of describing non-cosmological variance) and simplicity. The likelihood ratio and related tests can be used to assess the candidate extensions.

V. SUMMARY

Instrumental systematics, residual foregrounds, and other excess variance produce bias in cosmological parameters. Experiments to detect inflationary gravitational waves must use a battery of tests to rule out biases that could lead to a false detection. Determination of cosmological parameters directly from pixel-space likelihood has shown promise as a method to self-consistently handle foregrounds and survey depth variations or masks, especially on the largest scales in the survey. This approach bypasses calculation of the band powers, which have had a vigorous history of tests for systematics. We have developed some pixel-space likelihood analogies to the cross power, noise modeling, and the difference map null test.

If excess variance modulates with time or instrumental setup, a difference map can be interpreted using a likelihood for a particular source of parameter bias. The two dimensional posterior distribution of a parameter and its bias is a convenient diagnostic. We show examples of this parameter bias plane for the tensor to scalar ratio r when the excess variance is parameterized in the likelihood, and where it is not (mis-specification).

The parametric nature of the likelihood requires additional care. Bias in the pixel-space likelihood is signal-dependent because the map weights contain signal covariance. Signal dependence of the weight produces parameter distributions

that are not consistent between the sum and difference maps. We propose a procedure for consistently weighting contamination. The method accomplishes the role of a null test: under the same assumptions as the sum map analysis, is there evidence for parameter bias caused by modulated contamination in the difference?

We recommend an iterative approach. In a first pass, the likelihood models all cosmological parameters and imperfectly known instrumental terms (e.g. detector noise). If a weighted difference null test fails, that information should be used either to construct a model of excess variance or to reprocess the data in a way that eliminates the systematic effect. If a left-right scan difference fails, compensation of time constants should be reassessed until that test passes. Temperature-to-polarization leakage results in a covariance matrix between the temperature and the polarization. In parallel with cosmological parameters, the pixel-space likelihood should include any contamination which has a well-defined model.

The likelihood of the difference map provides a parametric test for mis-specification of the covariance model by isolating components that vary across data subsets. A more general problem is assessing whether there is variance in the data, time-varying or not, that is not explained well by the model and could produce spurious B-modes. A parametric model can be tested against less parametric models that are sensitive to a wider range of variance structure. In the case of CMB polarization, the power spectrum is already an excellent example of this approach and has been used by WMAP [22] and Planck [38] to corroborate likelihood results on large scales. The power spectrum exposes statistically isotropic variance with ℓ -dependence different from the signal. At the next level, tests for isotropy [27, 37] are sensitive to residual galactic foregrounds.

ACKNOWLEDGMENTS

We acknowledge useful comments from Jo Dunkley and an anonymous reviewer.

Appendix A: Relation of the joint likelihood and cross power approaches

For a model with only B-modes, $\mathbf{Q} = \mathbf{f}\mathbf{r}; \mathbf{r}_b \mathbf{g}$ the joint covariance across subseason maps A and B is

$$\mathbf{C} = \begin{pmatrix} (r + r_b)\mathbf{S}_T + \mathbf{N}_A & r\mathbf{S}_T \\ r\mathbf{S}_T & (r + r_b)\mathbf{S}_T + \mathbf{N}_B \end{pmatrix} \quad (\text{A1})$$

As a first iteration of the Newton-Raphson (NR) approach in Eq. (3), take the case where $r = r_b = 0$. In this case,

$$\mathbf{F} = \begin{pmatrix} \mathbf{Q} & \mathbf{0} \\ \mathbf{0} & \mathbf{Q} \end{pmatrix} \quad \mathbf{Q} = \frac{1}{2} \text{Tr}[(\mathbf{N}_A^{-1} + \mathbf{N}_B^{-1})\mathbf{S}_T(\mathbf{N}_A^{-1} + \mathbf{N}_B^{-1})\mathbf{S}_T] \quad \mathbf{Q} = \frac{1}{2} \text{Tr}[\mathbf{N}_A^{-1}\mathbf{S}_T\mathbf{N}_A^{-1}\mathbf{S}_T + \mathbf{N}_B^{-1}\mathbf{S}_T\mathbf{N}_B^{-1}\mathbf{S}_T] \quad (\text{A2})$$

and the initial iteration for \mathbf{r} is

$$\mathbf{r} = \frac{\text{Tr}(\mathbf{x}\mathbf{x}^T \mathbf{Q})}{2 \text{Tr}(\mathbf{N}_A^{-1}\mathbf{S}_T\mathbf{N}_B^{-1}\mathbf{S}_T)} \quad \mathbf{Q} = \mathbf{C}^{-1} \begin{pmatrix} 0 & \mathbf{S}_T \\ \mathbf{S}_T & 0 \end{pmatrix} \mathbf{C}^{-1} \quad (\text{A3})$$

where we have identified the form of an optimal quadratic estimator \mathbf{Q} .

The noise bias term

$$\text{Tr}[CQ] = \text{Tr} \begin{bmatrix} 0 & S_T \\ S_T & 0 \end{bmatrix} \begin{bmatrix} N_A^{-1} & 0 \\ 0 & N_B^{-1} \end{bmatrix} = 0. \quad (\text{A4})$$

The signal estimator term

$$\text{Tr}[xx^T Q] = 2x_A^T N_A^{-1} S_T N_B^{-1} x_B. \quad (\text{A5})$$

If we take the estimator \hat{r} to be this first NR iteration, it has the form of a cross power

$$\hat{r} = \frac{x_A^T N_A^{-1} S_T N_B^{-1} x_B}{\text{Tr}(N_A^{-1} S_T N_B^{-1} S_T)}. \quad (\text{A6})$$

Appendix B: Parameter dependence of the bias

Section III A argues that the bias in r from contamination covariance \square is $\text{Tr}(C^{-1} S_T C^{-1} \square) = \text{Tr}(C^{-1} S_T C^{-1} S_T)$. As a toy model to understand the behavior analytically, take the contaminant $\square = \square_b^{-2} 1$ and noise $N = \square_n^{-2} 1$. This is equivalent to having uniform detector noise in excess of what is predicted. Expand the B-mode signal covariance $r S_T = r U \square U^T$. The inverse covariance according to the Woodbury inverse is

$$C^{-1} = (r S_T + N)^{-1} = N^{-1} \square N^{-1} U (r \square^{-1} \square^{-1} + U^T N^{-1} U)^{-1} U^T N^{-1} \quad (\text{B1})$$

$$= \square_N^{-2} U [1 \square \square_N^{-2} (r \square^{-1} \square^{-1} + \square_N^{-2} 1)^{-1}] U^T. \quad (\text{B2})$$

Taking $r \neq 0$, $C^{-1} = \square_N^{-2} 1$. This is just the uniform detector noise weight. The quadratic estimator part of the bias trace is

$$C^{-1} S_T C^{-1} = \square_N^{-4} U [1 \square \square_N^{-2} (r \square^{-1} \square^{-1} + \square_N^{-2} 1)^{-1}] \square [1 \square \square_N^{-2} (r \square^{-1} \square^{-1} + \square_N^{-2} 1)^{-1}] U^T = \square_N^{-4} U W U^T; \quad (\text{B3})$$

where we have identified the weighting term W as the combination of factors between U and U^T . The weight is implemented as a vector multiplication as long as the contaminant is diagonalized by the same vectors U as the signal, which in this case is possible because we chose contaminant $\square = \square_b^{-2} 1$. Here,

$$\text{Tr}_{\text{bias}} = \frac{\text{Tr}(U W U^T \square_b^{-2} 1)}{\text{Tr}(U W U^T U \square U^T)} = \square_b^{-2} \frac{w(r)^T 1}{w(r)^T \square}; \quad (\text{B4})$$

where we have used the cyclic property of the trace, used the orthonormality of the eigenvectors, and let \square and w be the diagonals of the signal eigenvalues \square and weight W . In contrast, if the contaminant has the same covariance structure as the signal, $\text{Tr}_{\text{bias}} = \square_b^{-2} w(r)^T \square^{-1} w(r)^T \square = \square_b^{-2}$ with no dependence on r . Each term of the product $w(r)^T \square^{-1}$ is the \square $S = (S + N)$ of the signal mode. Modes in the map that contribute less signal to noise are downweighted. The denominator $\text{Tr}(C^{-1} S_T C^{-1} S_T) / w(r)^T \square$ is the effective number of independent modes of information about r in the map, and it is also equal to the Fisher matrix.

Appendix C: Curvature of the likelihood under mis-specification

The general curvature of the log-likelihood is [50]

$$2L_{,ij} = \text{Tr} [\square C^{-1} C_{,i} C^{-1} C_{,j} + C^{-1} C_{,ij} + C^{-1} (C_{,i} C^{-1} C_{,j} + C_{,j} C^{-1} C_{,i}) C^{-1} D \quad (\text{C1})$$

$$+ C^{-1} (C_{,i} C^{-1} D_{,j} + C_{,j} C^{-1} D_{,i}) - C^{-1} (C_{,ij} C^{-1} D + D_{,ij})]; \quad (\text{C2})$$

where $D = (x \square \square)(x \square \square)^T$. In the case that the covariance model C is the same as the covariance of the data $\text{hDi} = C$, and there is no parameter dependence in the mean, one recovers the usual Fisher matrix (expectation of the log-curvature)

$$F_{ij} = \frac{1}{2} \text{Tr}(C^{-1} C_{,i} C^{-1} C_{,j}); \quad (\text{C3})$$

With some unknown contamination x_c , the expectation value of D is $C + x_c x_c^T$. Generally the maximum-likelihood estimate of the covariance is biased by the presence of this contamination, or C_{biased} . Neglecting parameter dependence of the mean ($D_{,i}$) and second derivatives of the covariance (applicable for the models here with $r S_T$) gives

$$2L_{,ij} = \text{Tr} [\square C^{-1} C_{,i} C^{-1} C_{,j} + C^{-1} (C_{,i} C^{-1} C_{,j} + C_{,j} C^{-1} C_{,i}) C^{-1} \text{hDi}]; \quad (\text{C4})$$

The first term is the ordinary Fisher matrix, but we must replace $C^{-1} \rightarrow C_{\text{biased}}^{-1}$. The second term contains $C^{-1} h D i \rightarrow C_{\text{biased}}^{-1} (C + x_c x_c^T)$. The width of the likelihood (as explored in MCMC) no longer has a clear interpretation because the model in the likelihood C_{biased} does not coincide with $(C_{\text{true}} + x_c x_c^T)$.

-
- [1] Aiola, S., Amico, G., Battaglia, P., et al. 2012, in Proc. SPIE, Vol. 8446, Ground-based and Airborne Instrumentation for Astronomy IV, 84467A
 - [2] Allison, R., Caucal, P., Calabrese, E., Dunkley, J., & Louis, T. 2015, Phys. Rev. D, 92, 123535
 - [3] BICEP2 and Keck Array Collaborations, Ade, P. A. R., Ahmed, Z., et al. 2015, ApJ, 811, 126
 - [4] Bond, J. R., Jaffe, A. H., & Knox, L. 1998, Phys. Rev. D, 57, 2117
 - [5] Bunn, E. F., Zaldarriaga, M., Tegmark, M., & de Oliveira-Costa, A. 2003, Phys. Rev. D, 67, 023501
 - [6] Byrd, R. H., Lu, P., Nocedal, J., & Zhu, C. 1995, SIAM Journal on Scientific Computing, 16, 1190
 - [7] Das, S., Louis, T., Nolta, M. R., et al. 2014, JCAP, 4, 014
 - [8] Efstathiou, G. 2004, MNRAS, 348, 885
 - [9] —. 2004, MNRAS, 349, 603
 - [10] —. 2006, MNRAS, 370, 343
 - [11] Errard, J., Feeney, S. M., Peiris, H. V., & Jaffe, A. H. 2016, J. Cosmology Astropart. Phys., 3, 052
 - [12] Essinger-Hileman, T., Ali, A., Amiri, M., et al. 2014, in Proc. SPIE, Vol. 9153, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII, 91531I
 - [13] Farhang, M., Bond, J. R., Doré, O., & Netterfield, C. B. 2013, ApJ, 771, 12
 - [14] Gjerløw, E., Colombo, L. P. L., Eriksen, H. K., et al. 2015, ApJS, 221, 5
 - [15] Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759
 - [16] Grain, J., Tristram, M., & Stompor, R. 2009, Phys. Rev. D, 79, 123515
 - [17] Groeneboom, N. E., Eriksen, H. K., Gorski, K., et al. 2009, ApJ, 702, L87
 - [18] Gruetjen, H. F., & Shellard, E. P. S. 2014, Phys. Rev. D, 89, 063008
 - [19] Gruppuso, A., de Rosa, A., Cabella, P., et al. 2009, MNRAS, 400, 463
 - [20] Hamimeche, S., & Lewis, A. 2009, Phys. Rev. D, 79, 083012
 - [21] Henderson, S. W., Allison, R., Austermann, J., et al. 2016, Journal of Low Temperature Physics, 184, 772
 - [22] Hinshaw, G., Weiland, J. L., Hill, R. S., et al. 2009, ApJS, 180, 225
 - [23] Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, ApJS, 208, 19
 - [24] Hivon, E., Górski, K. M., Netterfield, C. B., et al. 2002, ApJ, 567, 2
 - [25] Hu, W., & White, M. 1997, New A, 2, 323
 - [26] Kamionkowski, M., Kosowsky, A., & Stebbins, A. 1997, Phys. Rev. D, 55, 7368
 - [27] Kamionkowski, M., & Kovetz, E. D. 2014, Physical Review Letters, 113, 191303
 - [28] Katayama, N., & Komatsu, E. 2011, ApJ, 737, 78
 - [29] Kogut, A., Fixsen, D. J., Chuss, D. T., et al. 2011, J. Cosmology Astropart. Phys., 7, 025
 - [30] Kouveliotou, C., Agol, E., Batalha, N., et al. 2014, ArXiv e-prints, arXiv:1401.3741
 - [31] Lazear, J., Ade, P. A. R., Benford, D., et al. 2014, in Proc. SPIE, Vol. 9153, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII, 91531L
 - [32] López-Caniego, M., Rebolo, R., Aguiar, M., et al. 2014, ArXiv e-prints, arXiv:1401.4690
 - [33] Mangilli, A., Płazczynski, S., & Tristram, M. 2015, MNRAS, 453, 3174
 - [34] Matsumura, T., Akiba, Y., Borrill, J., et al. 2014, Journal of Low Temperature Physics, 176, 733
 - [35] Page, L., Hinshaw, G., Komatsu, E., et al. 2007, ApJS, 170, 335
 - [36] Peebles, P. J. E. 1973, ApJ, 185, 413
 - [37] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014, A&A, 571, A23
 - [38] Planck Collaboration, Aghanim, N., Arnaud, M., et al. 2015, ArXiv e-prints, arXiv:1507.02704
 - [39] Remazeilles, M., Dickinson, C., Eriksen, H. K., & Wehus, I. K. 2016, MNRAS, 458, 2032
 - [40] Rocha, G., Contaldi, C. R., Bond, J. R., & Górski, K. M. 2011, MNRAS, 414, 823
 - [41] Rybicki, G. B., & Press, W. H. 1992, ApJ, 398, 169
 - [42] Seljak, U. 1998, ApJ, 503, 492
 - [43] Seljak, U., & Zaldarriaga, M. 1997, Physical Review Letters, 78, 2054
 - [44] Story, K. T., Reichardt, C. L., Hou, Z., et al. 2013, ApJ, 779, 86
 - [45] Tajima, O., Choi, J., Hazumi, M., et al. 2012, in Proc. SPIE, Vol. 8452, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI, 84521M
 - [46] Taylor, A. N., & Kitching, T. D. 2010, MNRAS, 408, 865
 - [47] Tegmark, M. 1997, Phys. Rev. D, 55, 5895
 - [48] Tegmark, M., & de Oliveira-Costa, A. 2001, Phys. Rev. D, 64, 063001
 - [49] Tegmark, M., Hamilton, A. J. S., Strauss, M. A., Vogele, M. S., & Szalay, A. S. 1998, ApJ, 499, 555
 - [50] Tegmark, M., Taylor, A. N., & Heavens, A. F. 1997, ApJ, 480, 22
 - [51] Tristram, M., Macías-Pérez, J. F., Renault, C., & Santos, D. 2005, MNRAS, 358, 833
 - [52] Watts, D. J., Larson, D., Marriage, T. A., et al. 2015, ApJ, 814, 103

- [53] Zhu, C., Byrd, R. H., Lu, P., & Nocedal, J. 1997, ACM Transactions on Mathematical Software (TOMS), 23, 550